Critical Infrastructure Protection
Quality + Engineering
Forensics - Assurance - Analytics

To: NTIA AI Request for Comment
From: Greg Hutchins PE CERM – Q+E Principal Engineer
Margaux Hutchins – Q+E Product Manager
Subject: AI Accountability Request For Comment
Date: June 10, 2023

This letter is in response to NTIA AI Accountability Request For Comment.

A little context as to the context of our responses:

- Quality Plus Engineering (Q+E) is a professional engineering, risk assurance company.
- Q+E was certified under the Safety Act for Critical Infrastructure Protection: Forensics, Assurance, Analytics® to conduct risk assurance and audits of critical infrastructure.
- Q+E has written best-selling books on risk based assurance and risk management such as **Value Added Auditing** (4 edition), **ISO 31000: ERM**, etc.
- Q+E has been involved in Lisp and rules based systems for more than 30 years.

Thank you for the opportunity to provide responses.

503.233.1012 or 800.COMPETE
Quality + Engineering
GregH@800Compete.com

AI Accountability Objectives

1. What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments? Responses could address the following:

a. What kinds of topics should AI accountability mechanisms cover? How should they be scoped?

The focus of our comments is on AI accountability, assurance, and trust of high risk, public-facing, decision making AI.

Our comments also focus on answering three questions regarding AI assurance:

1. What criteria will be used to evaluate AI adherence or compliance?
2. How will accountability, assurance and audits be conducted?
3. Who will conduct the audits?

The mechanisms should be doable and explainable. So, the assessments should be scoped narrowly focused on measurable attributes.

Accountability implies AI regulatory framework, risk controls, and even a risk taxonomy are already in place, commonly understood, and effectively deployed.

b. What are assessments or internal audits most useful for? What are external assessments or audits most useful for?

Assurance assessments are contextualized based on reporting requirements, level and type of assurance, and purpose. Internal audits and assessments provide assurance based on risk appetite and assurance. External audits share risk from auditee to auditor. External audits provide independent and objective assurance.

Three types of AI risk based, conformity assessments can be conducted: 1. First party; 2. Second party; and 3. third party. EU AIA endorses conformity assessment.

c. An audit or assessment may be used to verify a claim, verify compliance with legal standards, or assure compliance with non-binding trustworthy AI goals. Do these differences impact how audits or assessments are structured, credentialed, or communicated?

Yes, they change the risk – cost - benefit of the assurance assessment in terms of managing, planning, conducting, and reporting audits. Level and type of assurance will vary based on needs of parties and real/perceived AI risks.

Each of the above questions has 100's of managing, planning, conducting, and reporting questions.

d. Should AI audits or assessments be folded into other accountability mechanisms that focus on such goals as human rights, privacy protection, security, and diversity, equity, inclusion, and access? Are there benchmarks for these other accountability mechanisms that should inform AI accountability measures?

No. Audit goals such as ESG, human rights, equity, and inclusion are difficult to plan, conduct and assure. Why? They are variable, open to interpretation, lack of standards, and difficult to measure. If these questions can't be delineated, then the legislators and courts will spend many years parsing them.

For example, social goals are hugely important, but difficult to statistically define in terms of reasonableness, acceptable risk, acceptable bias, data consistency, accountability tracing, levels of confidence, etc.

e. Can AI accountability practices have meaningful impact in the absence of legal standards and enforceable risk thresholds? What is the role for courts, legislatures, and rulemaking bodies?

Legal standards, risk acceptance thresholds with confidence levels, and risk acceptance levels need to be developed that are actionable and meaningful. If not regulated properly, AI apps will be developed and platformed outside of U.S. much like crypto and other technologies.

Role of legislatures and courts will be to develop reasonable AI guard rails and assure reasonable deployment of high risk, public facing, and decision making AI.

2. Is the value of certifications, audits, and assessments mostly to promote trust for external stakeholders or is it to change internal processes? How might the answer influence policy design?

Policy design is predicated on many factors. AI assurance needs to be architected, designed, and deployed based on AI context. High risk, public facing, decision making AI is regulated and assured differently than low risk AI embedded in a product.

Yes. Assurance mechanisms provide requisite level of trust through risk transfer. Internal and external stakeholders will be beneficiaries.

Policy architecture, design and deployment should balance technical and social impacts.

3. AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there tradeoffs among these goals? To what extent can these inquiries be conducted by a single team or instrument?

a. The AI system does not substantially contribute to harmful discrimination against people.

This question needs to be parsed. Define 'substantially.' Define 'substantially contribute.' Define 'harmful discrimination.' These are vague terms that need to be auditable and operationalizable.

Unless these can be explicitly defined, assuring and accounting for AI systems will be difficult if not impossible.

An example may illustrate this. There is no absolute assurance. Automated or autonomous decision making is not free of bias. The public policy test will be to determine what is politically acceptable and reasonable.

b. The AI system does not substantially contribute to harmful misinformation, disinformation, and other forms of distortion and content-related harms.

Human oversight is essential throughout the AI development process. However, generative AI can hallucinate and provide output that is non causal or non-correlative making AI accountability difficult to assure.

Audits imply there is a causal or correlative relationship between inputs and outputs. But, what happens if the AI process is so opaque that audits cannot be conducted due to AI hallucinations, lack of explainability, etc.

Statement needs to be updated. Foreign actors and individual hackers are already using AI for asymmetric warfare. This is now a given.

AI has geo-political and geo-economic implications. AI can be developed by state actors, hackers, non-governmental institutions, and many others. AI has a low barrier of entry. AI is now global. Many accountability and assurance questions arise. How will the US or EU monitor and enforce global AI?

c. The AI system protects privacy.

The statement needs to be parsed.

This question seems moot. It was reported this week that DNA can be pulled from atmosphere, soil, water and snow. These can then be sequenced.

d. The AI system is legal, safe, and effective.

To paraphrase Rumsfeld, generative AI results in unintended consequences. There are many unknown unknowns regarding generative AI architecture, design, deployment, and assurance.

The statement needs to be parsed and framed. Define terms such as 'safe' and 'effective' in relation to different types of AI systems and decision making risks.

e. There has been adequate transparency and explanation to affected people about the uses, capabilities, and limitations of the AI system.

AI developers will be asked and/or be required to disclose proprietary and confidential AI information. AI regulation may require risk controls, disclosure, cost, transparency and other requirements. Stakeholders such as impacted parties and developers will need to know the AI guard rails and rules of development.

Developers may not want to disclose source code, capabilities, and risk (limitations) of their AI system unless they are indemnified. If this is legislated, then they will seek work arounds since AI is a global market.

f. There are adequate human alternatives, consideration, and fallbacks in place throughout the AI system lifecycle.

There are many human tools that can be used. There are gated reviews. Statistical analyses. PLC controls. Data set reviews. Corrective/Preventive actions. Compliance checks. Risk assessments. Simulations. Etc.

g. There has been adequate consultation with, and there are adequate means of contestation and redress for, individuals affected by AI system outputs.

This needs to be contextualized, parsed, and framed to define contestation and litigation thresholds. There is no such thing as absolute assurance. Define guard rails. Define 'adequate.' Define 'reasonable.'

Technical AI can be assured if standards are developed. Social AI systems are more difficult to assure.

h. There is adequate management within the entity deploying the AI system such that there are clear lines of responsibility and appropriate skillsets.

Clear lines of authority and accountability can be defined and deployed for AI development.

4. Can AI accountability mechanisms effectively deal with systemic and/or collective risks of harm, for example, with respect to worker and workplace health and safety, the health and safety of marginalized communities, the democratic process, human autonomy, or emergent risks?

Each of these stakeholders will be impacted by automated decision making and problem solving. Guard rails and terms need to be strictly defined as well as what is deemed legally acceptable and reasonable.

5. Given the likely integration of generative AI tools such as large language models (*e.g.*, ChatGPT) or other general-purpose AI or foundational models into downstream products, how can AI accountability mechanisms inform people about how such tools are operating and/or whether the tools comply with standards for trustworthy AI?

Generative AI is already integrated into apps, IOT devices, and most if not all products. Safety or assurance methods such as AI licensing and testing of high risk, public facing, decision making systems, stickers (CE mark) for IOT products, and self certification are some examples.

EU AI Act, NIST AI RMF, ISO 42001, ISO 19011-2018, AI NSW etc. offer principles and guidelines for trustworthy AI. Trustworthy principles and guidelines need to be operationalized and assured. The 'who', 'what', and 'how' questions need to be defined.

6. The application of accountability measures (whether voluntary or regulatory) is more straightforward for some trustworthy AI goals than for others. With respect to which trustworthy AI goals are there existing requirements or standards? Are there any trustworthy AI goals that are not amenable to requirements or standards? How should accountability policies, whether governmental or non-governmental, treat these differences?

This is a dissertation question. How to assure generative AI systems for race, gender, ethnicity and other protected characteristics when the systems can be variable and hallucinogenic are yet to be understood.

Safety, bias, trust, and other social goals are very important. How they will be AI regulated and managed are risks to be addressed.

7. Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?

Too much regulation, costly, and unfit regulation will prevent trustworthy AI.

Define minimum or reasonable standards of 'due care' and compliance.

8. What are the best definitions of and relationships between AI accountability, assurance, assessments, audits, and other relevant terms?

Red Book, Yellow Book, ISO, IA, AICIPA, COSO, IEEE, and others offer generic assurance guidelines and glossaries.

ISO is developing ISO 42001, which are 'what to adhere to' audit requirements. Canada is testing this AI conformity assessment model. ISO, IEC, and other standards making bodies are still several years away with a full set of AI guidelines. Each standard making approach has benefits and challenges. For example, ISO AI standards are narrow in scope and may cost up to $4,000 to obtain full set of AI guidelines.

AI standards being developed are also guidelines. They are open to interpretation and variable deployment. The hallmark of AI accountability and assurance is consistency of design and deployment.

Existing Resources and Models

9. What AI accountability mechanisms are currently being used? Are the accountability frameworks of certain sectors, industries, or market participants especially mature as compared to others? Which industry, civil society, or governmental accountability instruments, guidelines, or policies are most appropriate for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?

If AI is an existential risk, then global mutual recognition agreements can be developed such as air travel, nuclear power and similar high risk technologies.

Every regulatory agency has risk based enforcement regulations and requirements. US has FDA, FAA, and other regulatory inspection and testing schemes. Automated AI risk based, decision making regulation will need to be contextualized.

Final thought: Energy in the U.S. is regulated through DOE, NRC, FERC, NERC, etc. These schemes and frameworks work well. AI is different. AI may requires a multinational approach such as nuclear power proliferation agreements, testing, and monitoring.

10. What are the best definitions of terms frequently used in accountability policies, such as fair, safe, effective, transparent, and trustworthy? Where can terms have the same meanings across sectors and jurisdictions? Where do terms necessarily have different meanings depending on the jurisdiction, sector, or use case?

ISO, IEC, IEEE, and other standards making bodies have developed AI risk taxonomies. However, they need to be tailored to the application and stakeholder requirements.

Social terms such as fair, safe, effective, etc. need to be contextualized and operationalized for accountability and assurance.

11. What lessons can be learned from accountability processes and policies in cybersecurity, privacy, finance, or other areas?

There are a few lessons learned regarding accountability and assurance:

Regulation does not seem to address uncertainty and risk well. See recent banking failures.

Cybersecurity is notoriously porous. NIST 800's are guidelines. They are open to interpretation

Financial based auditing is almost 100 years old. PCAOB inspections in the U.S. fail 1/5 of public accounting firm audits.

'Check the box' approach to assurance will not work for high risk, generative AI.

Critical infrastructure with embedded AI systems are high risk such as nuclear power and water management. These high risk, technical systems will require engineering risk assurance rather than accounting/auditing assurance.

12. What aspects of the United States and global financial assurance systems provide useful and achievable models for AI accountability?

See above. Financial assurance is concentrated in 4 global firms. Financial statements are porous in many parts of the world.

IFRS and ISSB standards setting for ESG is a cautionary tale.

AI assurance and accountability requires a new approach to regulation, governance, accountability, and assurance.

13. What aspects of human rights and/or industry Environmental, Social, and Governance (ESG) assurance systems can and should be adopted for AI accountability?

IFRS and ISSB are setting standards for ESG. They are largely transparent. However, many questions still need to be addressed:

- Who conducts the audits?
- What are auditors checking adherence to?
- How are they conducting the audits?

Each of the above who, what and how questions have manifold sub issues that have yet to be decided.

14. Which non-U.S. or U.S. (federal, state, or local) laws and regulations already requiring an AI audit, assessment, or other accountability mechanism are most useful and why? Which are least useful and why?

NY Law 414 is a good start. Colorado and other states are developing similar initiatives. Most are struggling with the 'how', 'who', and 'what' questions.

EU, NSW (Aus), Canada, China, and many countries have national centric approaches.

Assurance and accountability will be tailored to the context and use cases. For example, Q+E developed the following for Critical Infrastructure Protection: Forensics, Assurance, Analytics® risk assurance:

- **Analytical.** Q+E engineers and scientists conduct risk and vulnerability assessments following Q+E protocols evaluating business continuity, cyber security, and physical

security systems against IEEE, PPD, NFPA, ISA, PMI, FISMA, ISO, NIST, CARVER, COSO, NERC, API, AGA, RAMCAP, RAM-T, FERC/NERC guidance, and ASIS standards.

- **Assurance.** Q+E offers the client three levels of assurance:
  - **Compliance.** Q+E conducts a compliance audit against appropriate standards and guidance.
  - **Assurance with opinion.** Q+E issues an opinion based on the results of a governance, risk, and compliance (GRC) audit or ERM risk-controls assessment.
  - **Assurance with insurance coverage.** Q+E conducts an audit and provides the requisite level of due diligence for the auditee to be covered.
- **Forensic.** Q+E provides the above levels of assurance as well as supplies a letter to the regulatory authority averring compliance that criteria have been met and internal controls are effective based on organization's risk appetite.

The above model is based on the premise that higher risk requires higher assurance.

Accountability Subjects

15. The AI value or supply chain is complex, often involving open source and proprietary products and downstream applications that are quite different from what AI system developers may initially have contemplated. Moreover, training data for AI systems may be acquired from multiple sources, including from the customer using the technology. Problems in AI systems may arise downstream at the deployment or customization stage or upstream during model development and data training.

a. Where in the value chain should accountability efforts focus?

Traceability of global software two or further tiers removed is difficult to assure. Blockchain and other technologies may help. However, AI verifiability, transparency, traceability, and ownership will need to be defined.

AI is bundled into many IOT and OEM products. Risks will emerge at each step of the product life cycle and sourcing lifecycle. OEM will own brand and product. Software developed offshore and 3 tiers away will be difficult to regulate. DOD, NIST and other agencies are developing software security requirements.

b. How can accountability efforts at different points in the value chain best be coordinated and communicated?

Very difficult to address as new geo-political and geo-economic axes develop.

For example, how does one audit and assure TikTok's AI algorithm?

c. How should vendors work with customers to perform AI audits and/or assessments? What is the role of audits or assessments in the commercial and/or public procurement process? Are there specific practices that would facilitate credible audits (*e.g.*, liability waivers)?

Risk based audits are expensive. They need to be scoped narrowly to first tier suppliers. They must have a value proposition to supplier stakeholders.

Indemnification and liability waivers are one of many possible inducements.

d. Since the effects and performance of an AI system will depend on the context in which it is deployed, how can accountability measures accommodate unknowns about ultimate downstream implementation?

Unknown unknowns was discussed previously.

16. The lifecycle of any given AI system or component also presents distinct junctures for assessment, audit, and other measures. For example, in the case of bias, it has been shown that "bias is prevalent in the assumptions about which data should be used, what AI models should be developed, where the AI system should be placed — or if AI is required at all." How should AI accountability mechanisms consider the AI lifecycle? Responses could address the following:

a. Should AI accountability mechanisms focus narrowly on the technical characteristics of a defined model and relevant data? Or should they feature other aspects of the socio-technical system, including the system in which the AI is embedded? When is the narrower scope better and when is the broader better? How can the scope and limitations of the accountability mechanism be effectively communicated to outside stakeholders?

Complex question. It is based on context, risk, and application. Part of this was discussed previously.

We recommend a narrow approach for AI technical risk assurance. Narrow approach is doable.

Tier AI apps based on public facing, decision making risk. For example, the higher the AI risk, higher the required assurance and accountability. High risk AI may require a professional engineering opinion or engineering attestation.

Social approach to accountability and assurance is much more difficult.

b. How should AI audits or assessments be timed? At what stage of design, development, and deployment should they take place to provide meaningful accountability?

Depending on AI risk profile, risk assurance can be cursory or throughout the product life cycle and deployment.

Any change to the AI model may also trigger a risk review. AI CAPA's can be developed as required.

c. How often should audits or assessments be conducted, and what are the factors that should inform this decision? How can entities operationalize the notion of continuous auditing and communicate the results?

Assurance should be based on the risks posed by the AI system. AI development can be monitored throughout the development and use lifecycle. Continuous auditing can be automated. Also, any material change to the AI system may trigger a risk review.

d. What specific language should be incorporated into governmental or non-governmental policies to secure the appropriate timing of audits or assessments?

Depends on context, stakeholder requirements, and risk factors. Language needs to be tailored to AI risk profile and use cases.

17. How should AI accountability measures be scoped (whether voluntary or mandatory) depending on the risk of the technology and/or of the deployment context? If so, how should risk be calculated and by whom?

See OECD guidelines and EU's AIA legislation.

18. Should AI systems be released with quality assurance certifications, especially if they are higher risk?

Quality assurance certifications can be released for low risk AI systems. Higher risk AI systems will require more robust assurance. Quality assurance such as ISO 9001 and similar quality management systems are low level assurance mechanisms.

Quality assurance specifications and ISO may be used for low level AI risk assurance, specifically assessment and compliance. However, they do not provide requisite assurance for medium to high risk AI. Why?

The essence of quality assurance is consistency in the management, planning, conducting and reporting audits. This is currently an ISO problem. Certification Body auditors often do not comply with ISO 19011-2018 requirements that mandate how to manage, plan, conduct, and report risk based audits. ISO auditors often lack technical and assurance proficiency. Forty hours of education is the only requirement to become an ISO lead assessor. Lack of audit consistency is another major challenge with ISO audits and conformity assessment.

AI self-certification is another option for low risk AI development.

19. As governments at all levels increase their use of AI systems, what should the public expect in terms of audits and assessments of AI systems deployed as part of public programs? Should the accountability practices for AI systems deployed in the public sector differ from those used for private sector AI? How can government procurement practices help create a productive AI accountability ecosystem?

Users expect trust. Trust assurance varies by user and use case. Public facing, risk based AI will require higher levels of assurance, transparency and trust.

Accountability Inputs and Transparency

20. What sorts of records (*e.g.*, logs, versions, model selection, data selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability? How long should this documentation be retained? Are there design principles (including technical design) for AI systems that would foster accountability-by-design?

Questions need to be parsed. Too broad.

Developers will also require trust mechanisms. How will code reviews, source code, and many other proprietary, confidential, classified, and sensitive information be protected if they have to be disclosed and assured.

If AI regulation is perceived as onerous and costly, then developers will do work arounds. If the risks are too high for AI development, executives will plead deniability, outsource development, deploy systems offshore, or do work arounds.

21. What are the obstacles to the flow of information necessary for AI accountability either within an organization or to outside examiners? What policies might ease researcher and other third-party access to inputs necessary to conduct AI audits or assessments?

See previous response.

22. How should the accountability process address data quality and data voids of different kinds? For example, in the context of automated employment decision tools, there may be no historical data available for assessing the performance of a newly deployed, custom-built tool. For a tool deployed by other firms, there may be data a vendor has access to, but the audited firm itself lacks. In some cases, the vendor itself may have intentionally limited its own data collection and access for privacy and security purposes. How should AI accountability requirements or practices deal with these data issues? What should be the roles of government, civil society, and academia in providing useful data sets (synthetic or otherwise) to fill gaps and create equitable access to data?

Refer to NY Law 144 and EU AIA for challenges and possible solutions.

23. How should AI accountability "products" (*e.g.*, audit results) be communicated to different stakeholders? Should there be standardized reporting within a sector and/or across sectors? How should the translational work of communicating AI accountability results to affected people and communities be done and supported?

Very critical social-technical question. A few observations:

Standardized reporting within a sector would communicate comparable risks, assurance and results. Cross sector reporting would be difficult due to variability of standards, use cases, data, regulations, risks, etc.

24. What are the most significant barriers to effective AI accountability in the private sector, including barriers to independent AI audits, whether cooperative or adversarial? What are the best strategies and interventions to overcome these barriers?

Barriers include: liability protections, IP protections, AI risk taxonomy, reasonableness, risk appetite/tolerance, etc. See previous responses.

25. Is the lack of a general federal data protection or privacy law a barrier to effective AI accountability?

Yes. EU is five years ahead of U.S. EU has developed PII protections.

Technology advances are making this question moot. For example, if DNA can be taken from atmosphere and sequenced, then technology may preempt many privacy and data protections.

26. Is the lack of a federal law focused on AI systems a barrier to effective AI accountability?

In general, yes. Federal law urgently needs to develop doable and assurable guard rails and guidelines.

Again, this question needs to be parsed. What type of federal law? What type of AI systems? What is effective AI accountability? If simple AI concepts can't be defined, AI can't be regulated and managed.

27. What is the role of intellectual property rights, terms of service, contractual obligations, or other legal entitlements in fostering or impeding a robust AI accountability ecosystem? For example, do nondisclosure agreements or trade secret protections are the possible consequences of AI accountability requirements that might impose significant costs on regulated entities? Are there ways to reduce these costs? What are the best ways to consider costs in relation to benefits?

Yes. Risk based auditing with an opinion is expensive. What are the level, type, and nature of the assurance deliverables. It should be applied to high risk, public facing, decision making AI apps.

Many questions remain for small business and developers. What is the inducement for the AI auditee to reveal proprietary and confidential information? Auditee will require NDA's and other protections.

USPTO is struggling with these questions currently as well as most U.S. government departments and agencies.

29. How does the dearth of measurable standards or benchmarks impact the uptake of audits and assessments?

Assurance requires: 1. 'What is' standards to assure against; 2. 'How to' audit guidelines to assure generative systems to a reasonable level; 3. 'Who' requirements specifying knowledge, skills, and abilities of auditors and assessors.

Without these, the AI audit scheme will be variable and lack consistency, which is the hallmark of good audits.

AI Accountability Policies

30. What role should government policy have, if any, in the AI accountability ecosystem?

Develop understandable and applicable rules for managing, planning, conducting, and reporting risk based AI audits to a requisite level of assurance.

a. Should AI accountability policies and/or regulation be sectoral or horizontal, or some combination of the two?

Assurance should be contextual and sectoral. Each sector has its own stakeholders, risk requirements and objectives.

b. Should AI accountability regulation, if any, focus on inputs to audits or assessments (*e.g.*, documentation, data management, testing and validation), on increasing access to AI systems for auditors and researchers, on mandating accountability measures, and/or on some other aspect of the accountability ecosystem?

All of the above.

c. If a federal law focused on AI systems is desirable, what provisions would be particularly important to include? Which agency or agencies should be responsible for enforcing such a law, and what resources would they need to be successful?

This is a sector specific question. Government departments and agencies should enforce AI problem solving and decision making within their regulatory purview and knowledge domain.

d. What accountability practices should government (at any level) itself mandate for the AI systems the government uses?

Public facing, high risk, decision making AI should be tailored to the required level of reasonable assurance.

31. What specific activities should government fund to advance a strong AI accountability ecosystem?

AI is potentially as disruptive as nuclear energy. Many are saying that AI is an existential threat. If this is the case, then U.S. government investment and oversight should be comparable to the Manhattan Project. Global agreements should be in place for high risk decision making such as public safety, environment, nuclear war, etc.

32. What kinds of incentives should government explore to promote the use of AI accountability measures?

Several NTIA questions seem punitive and costly.

Incentives and positive reinforcement would be preferred mechanism to induce compliance.

33. How can government work with the private sector to incentivize the best documentation practices?

AI Pandora's Box is already open. Global regulators are chasing the technology, which is moving much faster than regulation.

Government should understand AI risk based, problem solving and risk based, decision making.

Government should look at NFT and block chain for lessons learned. In the absence of regulation, companies and AI developers will develop best commercial practices.

34. Is it important that there be uniformity of AI accountability requirements and/or practices across the United States? Across global jurisdictions? If so, is it important only within a sector or across sectors? What is the best way to achieve it? Alternatively, is harmonization or interoperability sufficient and what is the best way to achieve that?

AI will be embedded in all products and integrated into all services. AI will influence public problem solving and decision making. Public policy architecture should be flexible and tailored to the sector.